

Claims

We claim:

1. For use with a database system having a workload comprising a set of queries that have been executed on the database, a method for selecting a set of partitioned physical database structures for access by the database system in executing queries comprising the steps of:

compiling a pool of partitioned candidate structures by:

for each query, determining potentially relevant structures and associating at least one partitioning method with each structure;

selecting potentially relevant structures with associated partitioning methods to add to the pool of partitioned candidate structures;

augmenting the pool of partitioned candidate structures by determining generalized partitioned structures that may be relevant over a set of queries in the workload and adding them to the pool of partitioned candidate structures;

enumerating a set of partitioned physical structures from the pool of partitioned candidate structures.

2. The method of claim 1 comprising the step of examining the workload to form a set of constraints on structures that may be added to the pool of partitioned candidate structures.

3. The method of claim 2 wherein the set of constraints is a set of column-subsets on which structures can be partitioned.

4. The method of claim 3 wherein the set of column-subsets is generated by evaluating a total cost of all queries in the workload that reference a given column-set and selecting column-sets that have a relatively high total cost of queries.

5. The method of claim 1 wherein the step of associating at least one partitioning method with each potentially relevant structure is performed by evaluating the query and associating a range partitioning method with the potentially relevant structure if the query comprises a range selection predicate on a single column.

6. The method of claim 5 wherein the range partitioning method is specified as the single column in the range selection predicate and an ordered sequence of all boundary values of ranges over the single column.

7. The method of claim 1 wherein the step of associating at least one partitioning method with each potentially relevant structure is performed by associating a hash partitioning method with the potentially relevant structure.

8. The method of claim 7 wherein the hash partitioning method is specified by a set of column types and a number of partitions.

9. The method of claim 8 wherein the number of partitions is calculated by iteratively evaluating the cost of executing the query with numbers of partitions that range between upper and lower partition number limits and selecting the number of partitions that has the lowest cost.

10. The method of claim 9 wherein the upper limit is a number of distinct values in a column-subset being partitioned.

11. The method of claim 9 wherein the lower limit is a number of processors in the database system.

12. The method of claim 9 wherein the lower limit is one.
13. The method of claim 9 comprising the step of rejecting any number of partitions that results in a partition that exceeds an amount of available memory.
14. The method of claim 9 comprising the step of rejecting any number of partitions that results in a number of partitions that exceeds a preset partition number limit.
15. The method of claim 9 wherein a plurality of potentially relevant structures are joined for the query and wherein the hash partitioning method associated with each of the potentially relevant structures comprises an identical number of partitions.
16. The method of claim 15 wherein the upper partition number limit is the minimum of the following values: the maximum number of distinct values in one of the plurality of potentially relevant structures, the combined size of the plurality of potentially relevant structures divided by available memory, or a maximum number of partitions allowed by the database system.
17. The method of claim 1 wherein the step of selecting potentially relevant structures with associated partitioning methods is performed by selecting a set of potentially relevant structures that returns a lowest optimizer estimated cost for the query.
18. The method of claim 1 wherein the step of determining generalized partitioned structures is performed by merging partitioned structures in the pool of partitioned candidate structures.
19. The method of claim 18 wherein the step of merging partitioned structures is performed by recursively pair wise merging all the partitioned structures in the pool,

selecting a merged structure that provides a highest cost benefit with respect to the workload, adding the selected merged structure to the set of partitioned candidate structures, removing the partitioned structures that were merged to form the selected merged structure from the pool, and returning to the pair wise merging step.

20. The method of claim 18 wherein the step of merging partitioned structures is performed by associating at least one partitioning method with each merged partitioned structure.

21. The method of claim 20 wherein the partitioned candidate structures being merged all have range partitioning methods and wherein the partitioning method associated with the merged partitioned structure is determined by determining the range partitioning method for the merged structure that results in the lowest total cost of evaluating all queries to which the merged partitioned structure is relevant.

22. The method of claim 21 wherein the cost of evaluating all queries is computed by: estimating a cost of scanning a subset of partitions required to answer each query based on a size of partitions being scanned and assigning a fixed cost for accessing any partition in answering the query to accumulate a total cost for each query.

23. The method of claim 20 wherein the step of associating a partitioning method with a merged structure is performed by associating a partitioning method identical to one associated with another structure that is relevant to a query that the merged structure is relevant to.

24. The method of claim 1 wherein a user may impose a constraint that any potentially relevant structure must have a partitioning method associated with it that is identical to a partitioning method of the table that the structure references.

25. The method of claim 20 wherein the partitioned candidate structures being merged all have range partitioning methods and wherein the partitioning method associated with the merged partitioned structure is determined by selecting a range partition method based on one of the queries in the workload.

26. The method of claim 8 wherein the number of partitions is calculated by iteratively evaluating numbers of partitions that are multiples of a number of processors on the database system and determining a minimum number of partitions that results in a partition size less than or equal to an amount of memory allocated for partition storage.

27. For use with a database system having a workload comprising a set of queries that have been executed on the database, a computer readable medium having computer executable steps stored thereon for performing method steps for selecting a set of partitioned physical database structures for access by the database system in executing queries, the method steps comprising:

compiling a pool of partitioned candidate structures by:

for each query, determining potentially relevant structures and associating at least one partitioning method with each structure;

selecting potentially relevant structures with associated partitioning methods to add to the pool of partitioned candidate structures;

augmenting the pool of partitioned candidate structures by determining generalized partitioned structures that may be relevant over a set of queries in the workload and adding them to the pool of partitioned candidate structures;

enumerating a set of partitioned physical structures from the pool of partitioned candidate structures.

28. The computer readable medium of claim 27 comprising the step of examining the workload to form a set of constraints on structures that may be added to the pool of partitioned candidate structures.

29. The computer readable medium of claim 28 wherein the set of constraints is a set of column-subsets on which structures can be partitioned.

30. The computer readable medium of claim 29 wherein the set of column-subsets is generated by evaluating a total cost of all queries in the workload that reference

a given column-set and selecting column-sets that have a relatively high total cost of queries.

31. The computer readable medium of claim 27 wherein the step of associating at least one partitioning method with each potentially relevant structure is performed by evaluating the query and associating a range partitioning method with the potentially relevant structure if the query comprises a range selection predicate on a single column.

32. The computer readable medium of claim 31 wherein the range partitioning method is specified as the single column in the range selection predicate and an ordered sequence of all boundary values of ranges over the single column.

33. The computer readable medium of claim 27 wherein the step of associating at least one partitioning method with each potentially relevant structure is performed by associating a hash partitioning method with the potentially relevant structure.

34. The computer readable medium of claim 33 wherein the hash partitioning method is specified by a set of column types and a number of partitions.

35. The computer readable medium of claim 34 wherein the number of partitions is calculated by iteratively evaluating the cost of executing the query with numbers of partitions that range between upper and lower partition number limits and selecting the number of partitions that has the lowest cost.

36. The computer readable medium of claim 35 wherein the upper limit is a number of distinct values in a column-subset being partitioned.

37. The computer readable medium of claim 35 wherein the lower limit is a number of processors in the database system.

38. The computer readable medium of claim 35 wherein the lower limit is one.

39. The computer readable medium of claim 35 comprising the step of rejecting any number of partitions that results in a partition that exceeds an amount of available memory.

40. The computer readable medium of claim 35 comprising the step of rejecting any number of partitions that results in a number of partitions that exceeds a preset partition number limit.

41. The computer readable medium of claim 35 wherein a plurality of potentially relevant structures are joined for the query and wherein the hash partitioning method associated with each of the potentially relevant structures comprises an identical number of partitions.

42. The computer readable medium of claim 41 wherein the upper partition number limit is the minimum of the following values: the maximum number of distinct values in one of the plurality of potentially relevant structures, the combined size of the plurality of potentially relevant structures divided by available memory, or a maximum number of partitions allowed by the database system.

43. The computer readable medium of claim 27 wherein the step of selecting potentially relevant structures with associated partitioning methods is performed by selecting a set of potentially relevant structures that returns a lowest optimizer estimated cost for the query.

44. The computer readable medium of claim 27 wherein the step of determining generalized partitioned structures is performed by merging partitioned structures in the pool of partitioned candidate structures.

45. The computer readable medium of claim 44 wherein the step of merging partitioned structures is performed by recursively pair wise merging all the partitioned structures in the pool, selecting a merged structure that provides a highest cost benefit with respect to the workload, adding the selected merged structure to the set of partitioned candidate structures, removing the partitioned structures that were merged to form the selected merged structure from the pool, and returning to the pair wise merging step.

46. The computer readable medium of claim 44 wherein the step of merging partitioned structures is performed by associating at least one partitioning method with each merged partitioned structure.

47. The computer readable medium of claim 46 wherein the partitioned candidate structures being merged all have range partitioning methods and wherein the partitioning method associated with the merged partitioned structure is determined by determining the range partitioning method for the merged structure that results in the lowest total cost of evaluating all queries to which the merged partitioned structure is relevant.

48. The computer readable medium of claim 47 wherein the cost of evaluating all queries is computed by: estimating a cost of scanning a subset of partitions required to answer each query based on a size of partitions being scanned and assigning a fixed

cost for accessing any partition in answering the query to accumulate a total cost for each query.

49. The computer readable medium of claim 46 wherein the step of associating a partitioning method with a merged structure is performed by associating a partitioning method identical to one associated with another structure that is relevant to a query that the merged structure is relevant to.

50. The computer readable medium of claim 27 wherein a user may impose a constraint that any potentially relevant structure must have a partitioning method associated with it that is identical to a partitioning method of the table that the structure references.

51. The computer readable medium of claim 46 wherein the partitioned candidate structures being merged all have range partitioning methods and wherein the partitioning method associated with the merged partitioned structure is determined by selecting a range partition method based on one of the queries in the workload.

52. The computer readable medium of claim 34 wherein the number of partitions is calculated by iteratively evaluating numbers of partitions that are multiples of a number of processors on the database system and determining a minimum number of partitions that results in a partition size less than or equal to an amount of memory allocated for partition storage.

53. For use with a database system having a workload comprising a set of queries that have been executed on the database, an apparatus for selecting a set of partitioned physical database structures for access by the database system in executing queries comprising:

a candidate accumulator that compiles a pool of partitioned candidate structures comprising:

a structure partitioner that, for each query, determines potentially relevant structures and associates at least one partitioning method with each structure;

a structure selector for selecting potentially relevant structures with associated partitioning methods to add to the pool of partitioned candidate structures;

a structure constructor for augmenting the pool of partitioned candidate structures by determining generalized partitioned structures that may be relevant over a set of queries in the workload and adding them to the pool of partitioned candidate structures;

a candidate set enumerator for enumerating a set of partitioned physical structures from the pool of partitioned candidate structures.

54. The apparatus of claim 53 comprising a workload analysis to for examining the workload to form a set of constraints on structures that may be added to the pool of partitioned candidate structures.

55. The apparatus of claim 54 wherein the set of constraints is a set of column-subsets on which structures can be partitioned.

56. The apparatus of claim 53 wherein the structure partitioner associates at least one partitioning method with each potentially relevant structure by evaluating the

query and associating a range partitioning method with the potentially relevant structure if the query comprises a range selection predicate on a single column.

57. The apparatus of claim 53 wherein the structure partitioner associates at least one partitioning method with each potentially relevant structure by associating a hash partitioning method with the potentially relevant structure.

58. The apparatus of claim 53 wherein the structure selector selects potentially relevant structures with associated partitioning methods by selecting a set of potentially relevant structures that returns a lowest optimizer estimated cost for the query.

59. The apparatus of claim 53 wherein the structure constructor merges partitioned structures in the pool of partitioned candidate structures to augment the pool of candidates.

60. For use with a database system having a workload comprising a set of queries that have been executed on the database, an apparatus for selecting a set of partitioned physical database structures for access by the database system in executing queries comprising:

means for compiling a pool of partitioned candidate structures comprising:

means for determining, for each query, potentially relevant structures;

means for associating at least one partitioning method with each determined structure;

means for selecting potentially relevant structures with associated partitioning methods to add to the pool of partitioned candidate structures;

means for augmenting the pool of partitioned candidate structures by determining generalized partitioned structures that may be relevant over a set of queries in the workload and adding them to the pool of partitioned candidate structures;

means for enumerating a set of partitioned physical structures from the pool of partitioned candidate structures.

61. The apparatus of claim 60 comprising means for examining the workload to form a set of constraints on structures that may be added to the pool of partitioned candidate structures.

62. The apparatus of claim 61 wherein the set of constraints is a set of column-subsets on which structures can be partitioned.

63. The apparatus of claim 60, wherein the means for associating at least one partitioning method associates at least one partitioning method with each potentially relevant structure by evaluating the query and associating a range partitioning method

with the potentially relevant structure if the query comprises a range selection predicate on a single column.

64. The apparatus of claim 60 wherein the means for associating at least one partitioning method associates at least one partitioning method with each potentially relevant structure by associating a hash partitioning method with the potentially relevant structure.

65. The apparatus of claim 60 wherein the means for selecting structures selects potentially relevant structures with associated partitioning methods by selecting a set of potentially relevant structures that returns a lowest optimizer estimated cost for the query.

66. The apparatus of claim 60 wherein the means for augmenting the pool of candidates merges partitioned structures in the pool of partitioned candidate structures to augment the pool of candidates.